

Ramanan 1-1

AF  
JW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

**Patent Application**

Applicant(s): K. Ramanan et al.  
Case: 1-1  
Serial No.: 09/393,949  
Filing Date: September 10, 1999  
Group: 2663  
Examiner: Soon D. Hyun

I hereby certify that this paper is being deposited on this date with the U.S. Postal Service as first class mail addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Signature: *K. Ramanan* Date: June 29, 2005

Title: Method and Apparatus for Scheduling Traffic to Meet Quality of Service Requirements in a Communication Network

TRANSMITTAL OF SUPPLEMENTAL APPEAL BRIEF

Mail Stop Appeal Brief - Patents  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Sir:

Submitted herewith are the following documents relating to the above-identified patent application:

- (1) Response to Office Action; and
- (2) Supplemental Appeal Brief.

There is no additional fee due in conjunction with the response. In the event of non-payment or improper payment of a required fee, the Commissioner is authorized to charge or to credit **Ryan, Mason & Lewis, LLP Deposit Account No. 50-0762** as required to correct the error.

Respectfully submitted,

*Joseph B. Ryan*

Date: June 29, 2005

Joseph B. Ryan  
Attorney for Applicant(s)  
Reg. No. 37,922  
Ryan, Mason & Lewis, LLP  
90 Forest Avenue  
Locust Valley, NY 11560  
(516) 759-7517



Ramanan 1-1

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

**Patent Application**

Applicant(s): K. Ramanan et al.  
Case: 1-1  
Serial No.: 09/393,949  
Filing Date: September 10, 1999  
Group: 2663  
Examiner: Soon D. Hyun

I hereby certify that this paper is being deposited on this date with the U.S. Postal Service as first class mail addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Signature: *V. Bencivenga* Date: June 29, 2005

Title: Method and Apparatus for Scheduling Traffic to Meet Quality  
of Service Requirements in a Communication Network

---

RESPONSE TO OFFICE ACTION

Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Sir:

In response to the Office Action dated March 29, 2005 in the above-referenced application, Applicants hereby request reinstatement of the appeal. A Supplemental Appeal Brief is submitted concurrently herewith.

Respectfully submitted,

Date: June 29, 2005

Joseph B. Ryan  
Attorney for Applicant(s)  
Reg. No. 37,922  
Ryan, Mason & Lewis, LLP  
90 Forest Avenue  
Locust Valley, NY 11560  
(516) 759-7517



Ramanan 1-1

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

**Patent Application**

Applicant(s): K. Ramanan et al.  
Case: 1-1  
Serial No.: 09/393,949  
Filing Date: September 10, 1999  
Group: 2663  
Examiner: Soon D. Hyun

I hereby certify that this paper is being deposited on this date with the U.S. Postal Service as first class mail addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Signature: *[Handwritten Signature]* Date: June 29, 2005

Title: Method and Apparatus for Scheduling Traffic to Meet Quality  
of Service Requirements in a Communication Network

---

**SUPPLEMENTAL APPEAL BRIEF**

Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Sir:

This Supplemental Appeal Brief is submitted in response to the Office Action dated March 29, 2005 in the above-referenced application, in which the Examiner reopened prosecution in response to the Appeal Brief filed January 17, 2005.

Applicants have submitted concurrently herewith a response to the March 29, 2005 Office Action, requesting reinstatement of the appeal.

**REAL PARTY IN INTEREST**

The present application is assigned to Lucent Technologies Inc., as evidenced by an assignment recorded September 10, 1999 in the U.S. Patent and Trademark Office at Reel 010248, Frame 0505. The assignee Lucent Technologies Inc. is the real party in interest.

### RELATED APPEALS AND INTERFERENCES

There are no known related appeals or interferences.

### STATUS OF CLAIMS

The present application was filed on September 10, 1999 with claims 1-28. Claims 1-28 are currently pending in the application. Claims 1, 14, 27 and 28 are the independent claims.

Claims 1-3, 14-16, 27 and 28 stand rejected under 35 U.S.C. §102(e). Claims 4-13 and 17-26 are indicated as containing allowable subject matter. Claims 1-3, 14-16, 27 and 28 are appealed.

### STATUS OF AMENDMENTS

There have been no amendments filed subsequent to the final rejection.

### SUMMARY OF CLAIMED SUBJECT MATTER

The present invention is directed to methods, apparatus and other arrangements for scheduling data packets for transmission over a communication link in a network. In accordance with the invention, delay measures are computed for a plurality of packets including at least one packet from each of a plurality of queues, and a given one of the plurality of packets is selected for transmission based at least in part on a comparison of weighted versions of computed delay measures, such that the selected packet is the packet having the largest weighted delay associated therewith.

An illustrative embodiment of the invention is implemented in a node 10 of a communication network as shown in FIG. 1. The node 10, also referred to as node A, receives a set of input data flows 11 from, e.g., other nodes or information sources of the network. The node 10 communicates with another node 12, also referred to as node B, via an A-B communication link 14.

The node 10 as shown in FIG. 1 has  $N$  input data flows, each of which is associated with a corresponding class of users and is applied to a corresponding packet queue 15-1, 15-2, . . . 15- $N$  within the node 10.

The node 10 further includes a link scheduler 16 which retrieves packets from the queues 15-1, 15-2, . . . 15- $N$  for transmission over the link 14. In accordance with the invention, the link scheduler 16 implements a Largest Weighted Delay First (LWDF) scheduling policy based on a set

of positive weights  $\alpha = \alpha_1, \alpha_2, \dots, \alpha_N$ . Generally speaking, the LWDF scheduling policy associated with the vector  $\alpha$  is a non-preemptive policy that always chooses for transmission the user class for which, e.g., the quantity  $W_i / \alpha_i$  is maximum, where  $W_i$  is the delay experienced by the head-of-the-queue packet of user class  $i$ , i.e., the delay experienced by the packet at the head of the queue 15- $i$  in node 10. Advantageously, the form of the scheduling algorithm in this illustrative embodiment is invariant to changes in stochastic input flow models. Additional details regarding the scheduling policy of the illustrative embodiment are shown in the flow diagram of FIG. 2, as described in the specification at page 7, lines 8-27.

#### GROUND OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1-3, 14-16, 27 and 28 are anticipated under 35 U.S.C. §102(e) by U.S. Patent No. 6,108,307 (hereinafter “McConnell”).

#### ARGUMENT

MPEP §2131 specifies that a given claim is anticipated “only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference,” citing Verdegaal Bros. v. Union Oil Co. of California, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). Moreover, MPEP §2131 indicates that the cited reference must show the “identical invention . . . in as complete detail as is contained in the . . . claim,” citing Richardson v. Suzuki Motor Co., 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). For the reasons identified below, Applicants submit that the Examiner has failed to establish anticipation of at least independent claims 1, 14, 27 and 28 by McConnell.

Each of independent claims 1, 14, 27 and 28 includes a limitation that generally specifies the selection of a given one of a plurality of packets for transmission based at least in part on a comparison of weighted versions of computed delay measures, such that the selected packet is the packet having the largest weighted delay associated therewith.

Also, it is important to note that the claims require that the delay measures are computed for a plurality of packets including at least one packet from each of a plurality of queues. Thus, the “computed delay measures” recited in the claims refer to delay measures computed for packets from different queues.

Accordingly, the claims call for computing delay measures for packets from different queues, and comparing weighted versions of the computed delay measures, such that the packet having the largest weighted delay is selected for transmission.

As noted above, an illustrative embodiment of the claimed invention as described in the specification at page 4, line 12 to page 5, line 28, provides an improved scheduling policy referred to as Largest Weighted Delay First (LWDF). One important advantage associated with this embodiment is that the LWDF scheduling policy is “invariant to changes in stochastic input flow models” (Specification, page 5, lines 26-28).

The McConnell reference fails to teach or suggest at least the above-noted limitations of claims 1, 14, 27 and 28, and thus also fails to provide the associated advantages, such as invariability to changes in stochastic input flow models.

The Examiner in formulating the §102(e) rejection relies on the frame processing operations associated with queues Q1 through Q4 in FIG. 3 of McConnell. Applicants respectfully submit that frames from these queues are not processed in a manner which anticipates the above-described limitations of the independent claims.

Applicants initially note that McConnell indicates that received frames are assigned to the queues Q1 through Q4 based on “predetermined priority levels.” This is apparent from, for example, the disclosure in column 5, lines 50-63, which provides as follows with emphasis supplied:

When a frame arrives, the frame processor reads the DLCI bits 30 in address field 15 of the frame message. The frame processor accesses a lookup table 47 located in memory space 48 in order to assign a predetermined priority level for a message having the DLCI of the message received by the frame processor 46.

As explained in greater detail below, the various predetermined priority levels found in lookup table 47 will each correspond to a respective first-in first-out (FIFO) queue 50, 52, 54, 56 of memory space 48. Preferably, each of these queues is of fixed length. The frame processor 46 places the received frame into the tail of one of the queues 50, 52, 54, 56 according to its associated priority level provided by lookup table 47.

Thus, McConnell teaches that the queues themselves are already prioritized. Additional disclosure regarding this predetermined prioritization is found in column 6, lines 16-30, and provides as follows with emphasis supplied:

As explained previously, the number of queues implemented in the frame service device 43 will reflect the number of service class priority levels to be handled by the congestion management scheme. The priority levels may be arbitrarily assigned from low to high, or may be predetermined to reflect or map to quality of service parameters associated with an ATM backbone network. In the preferred embodiment of the frame service device described above, four queues are implemented, which each correspond to a particular class of service associated with the virtual connections of the network node to which the frame processing apparatus 43 pertains. For instance, the queues 50, 52, 54, 56 may respectively represent high priority, medium priority, low priority and best effort priority levels, respectively.

Moreover, McConnell indicates that the predetermined priority levels are used, in conjunction with congestion severity status determined for each of the queues Q1 through Q4, to select a frame from a particular one of the queues Q1 through Q4. This is apparent from the disclosure in column 9, lines 12-20, which provides as follows with emphasis supplied:

At step 100, it is determined whether there exists more than one frame having the highest queue congestion severity. If not, at step 102 the frame server 58 dispatches the frame with the highest queue congestion severity. However, if there exists more than one frame with the highest queue congestion severity, this congestion severity tie will be resolved in favour of the highest priority level queue at step 104.

Applicants submit that the “congestion severity” measure determined for each of the queues in McConnell does not constitute a weighted version of a computed delay measure, as claimed. Instead, the congestion severity is determined by the use of depth congestion thresholds and age congestion thresholds, as described at column 7, line 10, to column 8, line 17, and without any use of weighting whatsoever. McConnell therefore fails to teach or suggest the claimed arrangements

involving computing delay measures for packets from different queues, and comparing weighted versions of the computed delay measures, such that the packet having the largest weighted delay is selected for transmission. In fact, McConnell could be said to teach away from these limitations, by explicitly disclosing a frame selection mechanism which is implemented without any type of weighting of computed delay measures.

The Examiner in a prior Office Action dated August 10, 2004 argues that the congestion severity measure constitutes a weighted version of the age severity measure, where age severity corresponds to delay and the weighting is by the depth severity measure (August 10, 2004 Office Action, page 3, first full paragraph). However, this interpretation is inconsistent with the explicit teachings of McConnell. For example, McConnell teaches in column 8, lines 10-17, that the age severity and depth severity are treated independently in establishing the congestion severity. It cannot be said that the depth severity is applied as a weighting to the age severity, to create a “weighted version” of the age severity, as would be required by the claim language. Instead, both depth severity and age severity have an entirely separate and independent influence on the congestion severity.

The Examiner in the August 10, 2004 Office Action, at page 4, first four paragraphs, again argues that the congestion severity measure constitutes a weighted version of the age severity measure, where age severity corresponds to delay and the weighting is by the depth severity measure. The Examiner relies on the teachings in McConnell at column 7, lines 10-15, and steps 72, 78, 88 and 102 of the flow diagram in FIG. 6. Applicants believe that the relied-upon portions of McConnell fail to support the anticipation argument of the Examiner. The McConnell reference at column 7, line 10, to column 8, line 17, teaches as follows regarding the relationship between the congestion severity measure, the depth severity measure, and the age severity measure, with emphasis supplied:

The congestion severity status of each queue is determined whenever the frame server 58 is ready to dispatch a frame from the head of the queues. In the preferred embodiment of the present invention, the congestion severity status is defined by two components, namely a depth severity component and an age severity component.



The congestion severity status for a queue is determined as the highest severity status pertaining to either of the depth severity component or the age severity component associated with that queue. Thus, if the depth severity component has a higher severity than the age severity component, the former defines the congestion severity status for the queue. However, if the depth severity component has a lower severity than the age severity component, the latter defines the congestion severity for the queue. If both the depth and age severity are the same, that severity is equated to the congestion severity component for the queue.

The depth severity component for a queue is determined by comparing the current depth of the queue with the three depth congestion thresholds applied to the queue. The greater the number of applied depth congestion thresholds which have been exceeded by the current measure of queue depth, the greater the depth severity. Where depth congestion thresholds have been applied as previously explained, a colour scheme as follows may be employed to denote depth severity:

Depth Severity	
Depth Status	
Green (Low)	The current queue depth is less than the MCT depth congestion level.
Yellow (Mid)	The current queue depth is greater than or equal to the MCT depth congestion level but less than the SCT depth congestion level
Red (High)	The current queue depth is greater than or equal to the SCT depth congestion level.

The age severity component for a queue is determined by comparing the current age of the message frame at the head of the particular queue, calculated according to its time stamp information previously described, with the three age congestion thresholds applied to the queue. The greater the number of applied age congestion thresholds which have been exceeded by the current measure of age for the message frame, the greater the age severity

for the queue. Where age congestion thresholds have been applied as previously explained, a colour scheme as follows may be employed to denote depth severity:

Age Severity	
Age Status	
Green (Low)	The current age of the message frame at the head of the queue is less than the MCT age congestion level.
Yellow (Mild)	The current age of the message frame at the head of the queue is greater than or equal to the MCT age congestion level but less than the SCT age congestion level.
Red (High)	The current age of the message frame at the head of the queue is greater than or equal to the SCT age congestion level.

Accordingly, in the preceding examples, the congestion severity status for a queue will result in a determination of Red (or High) if either or both of the depth severity and age severity is Red (or High); a determination of Yellow (or Mid) if neither depth severity or age severity is Red (or High) and either or both are Yellow (or Mid); and a determination of Green (or Low) if every depth severity and age severity is Green.

The above-quoted teachings from McConnell make it clear that the depth severity measure and the age severity measure are determined independently of one another. That is, the depth severity measure is determined by comparing the current depth of the queue with the three depth congestion thresholds applied to the queue, while the age severity measure is determined by comparing the current age of the message frame at the head of the particular queue with the three age congestion thresholds applied to the queue. Moreover, the depth severity measure and age severity measure exert entirely separate and independent influences on the congestion severity status. That is, if the depth severity measure and the age severity measure differ from one another, whichever measure is larger “defines the congestion severity status for the queue.” Thus, it is believed that the Examiner is incorrect in arguing that the congestion severity measure of McConnell somehow constitutes a weighted version of the age severity measure, where the weighting is by the depth

severity measure. See page 3, first paragraph, of the August 10, 2004 Office Action. The age severity measure is simply not “weighted” in any way by the depth severity measure in determining the congestion severity measure.

In the March 29, 2005 Office Action, at page 3, third paragraph, the Examiner relies on the column 7, lines 11-26, portion of the above-quoted teachings from McConnell regarding the relationship between the congestion severity measure, the depth severity measure, and the age severity measure. More specifically, the Examiner argues that because the congestion severity measure is determined using the depth severity measure and the age severity measure, a packet selected based on congestion severity status is selected based on a comparison of weighted versions of computed delay measures, such that the packet having the largest weighted delay is selected for transmission. Applicants respectfully disagree, for the reasons noted above. As described in detail previously herein, the congestion severity measure in McConnell cannot reasonably be viewed as a weighted version of any particular computed delay measure.

In view of the foregoing, Applicants respectfully submit that there is no teaching or suggestion in McConnell of the particular limitations of claims 1, 14, 27 and 28 relating to computing delay measures for a plurality of packets including at least one packet from each of a plurality of queues, and selecting a given one of the plurality of packets for transmission based at least in part on a comparison of weighted versions of the computed delay measures, such that the selected packet is the packet having the largest weighted delay associated therewith.


As indicated previously, the McConnell reference, by teaching to utilize an entirely different scheduling technique, actively teaches away from the present invention as claimed. More specifically, the McConnell reference at column 8, lines 31-41, that the frame server 58 “accepts the frame at the head of the particular queue which has the highest congestion severity status of all queues.” In the event of multiple queues having the highest congestion severity status, the predetermined priority levels of the queues are used to break the tie. Thus, McConnell teaches away from the claimed selection based on largest weighted delay.

Since McConnell fails to teach or suggest each and every element of claims 1, 14, 27 and 28, as would be required for an appropriate anticipation rejection, these claims are not anticipated by McConnell, and the §102(e) rejection should be withdrawn.

Dependent claims 2, 3, 15 and 16 are believed allowable at least by virtue of their dependence from their respective independent claims.

In view of the above, Applicants believe that claims 1-28 are in condition for allowance, and respectfully request the withdrawal of the §102(e) rejection.

Respectfully submitted,

A handwritten signature in black ink, reading "Joseph B. Ryan". The signature is fluid and cursive, with the first name "Joseph" and last name "Ryan" clearly legible.

Date: June 29, 2005

Joseph B. Ryan  
Attorney for Applicant(s)  
Reg. No. 37,922  
Ryan, Mason & Lewis, LLP  
90 Forest Avenue  
Locust Valley, NY 11560  
(516) 759-7517

## CLAIMS APPENDIX

1. A processor-implemented method of scheduling data packets for transmission over a communication link in a network, the method comprising the steps of:

computing delay measures for a plurality of packets including at least one packet from each of a plurality of queues; and

selecting a given one of the plurality of packets for transmission based at least in part on a comparison of weighted versions of the computed delay measures, such that the selected packet is the packet having the largest weighted delay associated therewith.

2. The method of claim 1 wherein the step of computing delay measures for a plurality of packets includes computing the delay measure for a given one of the packets as a difference between a current time and an arrival time of the given packet in a node of the communication network.

3. The method of claim 1 wherein the plurality of packets comprise  $N$  packets, each associated with a corresponding one of  $N$  data flows and located in a head position in a corresponding one of  $N$  queues.

4. The method of claim 3 wherein the selecting step includes selecting a packet from a head position of a  $j$ th one of the  $N$  queues, such that a function of a given weight  $\alpha_j$  and a given delay measure  $W_j$  is maximal among all such functions for  $j = 1, 2, \dots, N$ , where  $\alpha_j$  is one of a set of positive weights  $\alpha_1, \alpha_2, \dots, \alpha_N$ .

5. The method of claim 4 wherein the function of the given weight and the given delay measure is  $W_j/\alpha_j$ .

6. The method of claim 4 wherein the computing and selecting steps are configured to meet a quality of service requirement specified in terms of a deadline  $T_i$  and an allowed deadline violation probability  $\delta_i$ .

7. The method of claim 6 wherein each of the weights  $\alpha_i$  in the set of positive weights  $\alpha_1, \alpha_2, \dots, \alpha_N$  is given by  $\alpha_i = -T_i / \log \delta_i$ .

8. The method of claim 6 wherein the quality of service requirement is specified by:

$$P(W_i > T_i) \leq \delta_i \text{ for } i = 1, \dots, N.$$

9. The method of claim 3 wherein the selecting step includes selecting a packet from a head position of a  $j$ th one of the  $N$  queues, such that a function of a given weight  $\alpha_j$  and a given delay measure  $Q_j$  is maximal among all such functions for  $j = 1, 2, \dots, N$ , where  $\alpha_j$  is one of a set of positive weights  $\alpha_1, \alpha_2, \dots, \alpha_N$ , and  $Q_j$  is a queue length of the  $j$ th queue.

10. The method of claim 9 wherein the function of the given weight and the given delay measure is  $Q_j/\alpha_j$ .

11. The method of claim 10 wherein the computing and selecting steps are configured to meet a quality of service requirement specified in terms of a maximum queue length  $H_i$  and an allowed queue length violation probability  $\delta_i$ .

12. The method of claim 11 wherein each of the weights  $\alpha_i$  in the set of positive weights  $\alpha_1, \alpha_2, \dots, \alpha_N$  is given by  $\alpha_i = -H_i / \log \delta_i$ .

13. The method of claim 11 wherein the quality of service requirement is specified by:

$$P(Q_i > H_i) \leq \delta_i \text{ for } i = 1, \dots, N.$$

14. An apparatus for use in scheduling data packets for transmission over a communication link in a network, the apparatus comprising:

a set of queues, each of at least a subset of the queues including at least one packet;

and

a link scheduler having inputs coupled to corresponding outputs of the queues, the link scheduler being operative to compute delay measures for a plurality of packets including at least one packet from each of the at least a subset of the queues, and to select a given one of the plurality of packets for transmission based at least in part on a comparison of weighted versions of the

computed delay measures, such that the selected packet is the packet having the largest weighted delay associated therewith.

15. The apparatus of 14 wherein the link scheduler is further operative to compute delay measures for a plurality of packets includes computing the delay measure for a given one of the packets as a difference between a current time and an arrival time of the given packet in a node of the communication network.

16. The apparatus of claim 14 wherein the plurality of packets comprise  $N$  packets, each associated with a corresponding one of  $N$  data flows and located in a head position in a corresponding one of  $N$  queues.

17. The apparatus of claim 16 wherein the link scheduler is further operative to select a packet from a head position of a  $j$ th one of the  $N$  queues, such that a function of a given weight  $\alpha_j$  and a given delay measure  $W_j$  is maximal among all such functions for  $j = 1, 2, \dots, N$ , where  $\alpha_j$  is one of a set of positive weights  $\alpha_1, \alpha_2, \dots, \alpha_N$ .

18. The apparatus of claim 17 wherein the function of the given weight and the given delay measure is  $W_j/\alpha_j$ .



19. The apparatus of claim 17 wherein the link scheduler is configured to meet a quality of service requirement specified in terms of a deadline  $T_i$  and an allowed deadline violation probability  $\delta_i$ .

20. The apparatus of claim 19 wherein each of the weights  $\alpha_i$  in the set of positive weights  $\alpha_1, \alpha_2, \dots, \alpha_N$  is given by  $\alpha_i = -T_i / \log \delta_i$ .

21. The apparatus of claim 19 wherein the quality of service requirement is specified by:

$$P(W_i > T_i) \leq \delta_i \text{ for } i = 1, \dots, N.$$

22. The apparatus of claim 16 wherein the link scheduler is further operative to select a packet from a head position of a  $j$ th one of the  $N$  queues, such that a function of a given weight  $\alpha_j$  and a given delay measure  $Q_j$  is maximal among all such functions for  $j = 1, 2, \dots, N$ , where  $\alpha_j$  is one of a set of positive weights  $\alpha_1, \alpha_2, \dots, \alpha_N$ , and  $Q_j$  is a queue length of the  $j$ th queue.

23. The apparatus of claim 22 wherein the function of the given weight and the given delay measure is  $Q_j/\alpha_j$ .

24. The apparatus of claim 22 wherein the link scheduler is configured to meet a quality of service requirement specified in terms of a maximum queue length  $H_i$  and an allowed queue length violation probability  $\delta_i$ .

25. The apparatus of claim 19 wherein each of the weights  $\alpha_i$  in the set of positive weights  $\alpha_1, \alpha_2, \dots, \alpha_N$  is given by  $\alpha_i = -H_i / \log \delta_i$ .

26. The apparatus of claim 19 wherein the quality of service requirement is specified by:

$$P(Q_i > H_i) \leq \delta_i \text{ for } i = 1, \dots, N.$$

27. An apparatus for use in scheduling data packets for transmission over a communication link in a network, the apparatus comprising:

a memory implementing a set of queues, each of at least a subset of the queues including at least one packet; and

a processor coupled to the memory and operative to compute delay measures for a plurality of packets including at least one packet from each of the at least a subset of the queues, and to select a given one of the plurality of packets for transmission based at least in part on a comparison of weighted versions of the computed delay measures, such that the selected packet is the packet having the largest weighted delay associated therewith.

28. A machine-readable medium storing one or more software programs for scheduling data packets for transmission over a communication link in a network, wherein the one or more programs when executed by a processor include:

a step for computing delay measures for a plurality of packets including at least one packet from each of a plurality of queues; and

a step for selecting a given one of the plurality of packets for transmission based at least in part on a comparison of weighted versions of the computed delay measures, such that the selected packet is the packet having the largest weighted delay associated therewith.

EVIDENCE APPENDIX

None

RELATED PROCEEDINGS APPENDIX

None